U.S. DEPARTMENT OF ENERGY | Office of Science

Data Analysis through Advanced Scientific Computing Research at the Department of Energy

Presented at the 5th International Workshop on Data Analysis and Reduction for Big Scientific Data

Denver, CO

November 17, 2019

Bill Spotz, PhD

Office of Advanced Scientific Computing Research

Program Manager for Applied Mathematics

William.Spotz@science.doe.gov

Office of Science
By the numbers

Commissioned in 2018, Summit is the newest supercomputer at Oak Ridge National Laboratory. Compared to Titan, Summit achieves a factor of 8 increase in performance with one-fourth the number of nodes. Summit is an IBM AC922 system and links more than 27,000 NVIDIA Volta GPUs with more than 9,000 IBM Power9 CPUs. Its benchmark performance is 201 Petaflops. For some artificial intelligence applications, the selective use of lower-precision calculations can potentially quadruple Summit's performance to exascale levels.

SC delivers scientific discoveries and tools to transform our understanding of nature and advance the energy, economic, and national security of the U.S.

Research

- Provides about half of the U.S. Federal support for basic research in the physical sciences;
- Supports about 19,000 Ph.D. scientists, graduate students, engineers, and support staff at over 300 institutions and 10 DOE national laboratories;
- Maintains U.S. and world leadership in high-performance computing and computational sciences;
- Continues to be the major U.S. supporter of physics, chemistry, materials sciences, and biology for discovery and for energy sciences.

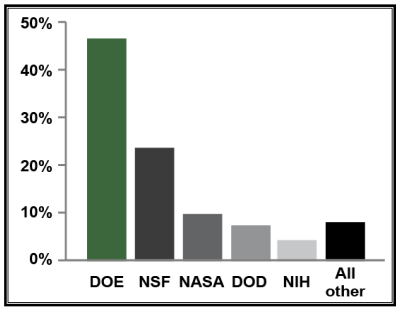

Support for basic research in the physical sciences by agency.

Source: *NSF Science and Engineering Indicators 2012*

Scientific User Facilities

- SC maintains the world's largest collection of scientific user facilities (aka research infrastructure) operated by a single organization in the world, used by more than 27,000 researchers each year.

Office of Science: User Facilities portfolio

- **Supercomputers,**
- **High intensity x-ray, neutron, and electron sources,**
- **Nanoscience facilities,**
- **Genomic sequencing facilities,**
- **Particle accelerators,**
- **Fusion/plasma physics facilities, and**
- **Atmospheric monitoring capabilities**

- **Open access; allocation determined through peer review of proposals**

- **Free for non-proprietary work published in the open literature**

- Full cost recovery for proprietary work

DOE Office of Science Laboratories

Berkeley, California
202 acres and 97 buildings
3,396 FTEs
950 students & postdocs
9,320 facility users
www.lbl.gov

Richland, Washington
346 acres and 19 buildings
4,344 FTEs
550 students & postdocs
1,733 facility users
www.pnnl.gov

Ames, Iowa
8 acres and 12 buildings
308 FTEs
158 students & postdocs
www.ameslab.gov

Batavia, Illinois
6,800 acres and 354 buildings
1,720 FTEs
55 students & postdocs
2,097 facility users
www.fnal.gov

Argonne, Illinois
1,517 acres and 100 buildings
3,460 FTEs
1,054 students & postdocs
6,547 facility users
www.anl.gov

Menlo Park, California
426 acres and 151 buildings
1,596 FTEs
213 students & postdocs
4,474 facility users
www.slac.stanford.edu

Oak Ridge, Tennessee
4,421 acres and 194 buildings
4,586 FTEs
1,080 students & postdocs
3,215 facility users
www.ornl.gov

Newport News, Virginia
169 acres and 72 buildings
729 FTEs
60 students & postdocs
1,261 facility users
www.jlab.org

Princeton, New Jersey
89 acres and 34 buildings
429 FTEs
54 students & postdocs
290 facility users
www.pppl.gov

Upton, New York
5,322 acres and 310 buildings
2,882 FTEs
642 students & postdocs
4,134 facility users
www.bnl.gov

Office of Science: Research portfolio

Advanced Scientific Computing Research	• **Delivering world leading computational and networking capabilities to extend the frontiers of science and technology**
Basic Energy Sciences	• **Understanding, predicting, and ultimately controlling matter and energy flow at the electronic, atomic, and molecular levels**
Biological and Environmental Research	• **Understanding complex biological, climatic, and environmental systems**
Fusion Energy Sciences	• **Building the scientific foundations for a fusion energy source**
High Energy Physics	• **Understanding how the universe works at its most fundamental level**
Nuclear Physics	• **Discovering, exploring, and understanding all forms of nuclear matter**

Overview of Advanced Scientific Computing Research

Computational and networking capabilities to extend the frontiers of science and technology

- **Basic Research**
  - Applied Mathematics research focused on addressing challenges of understanding increasingly complex natural and engineering system. Another focus is on improving the rigor and reliability of machine learning and artificial intelligence techniques.
  - Computer Science research enables computing and networking at extreme scales, and the understanding of extreme scale or complex data from both simulations and experiments. It addresses challenges posed by increased heterogeneity and complexity of computing systems, and the need to integrate simulation, data analysis, and other tasks into coherent workflows.
- **Computational Partnerships:** SciDAC Institutes and Science Application Partnerships connect ASCR basic research with domain scientists. New efforts include teams and co-design partnerships for Quantum Information Science and for data science at Scientific User Facilities.
- **High Performance Computing and Networking Facilities** operate with >90% availability.
- **Next Generation Computing**
  - Near Term: Exascale Computing Initiative (ECI) and Exascale Computing Project (ECP). The ECI activity is a joint ASCR/NNSA partnership to undertake, through ECP, the application, software and hardware R&D necessary to develop an exascale ecosystem and through the facilities, deploy at least one exascale-capable computer (10^{18} operations per second) in 2021.
  - Far Term: Research and Evaluation Partnerships explore future computing technologies, including quantum testbeds.
- **Workforce**: Computational Sciences Graduate Fellowship, Early Career Research Program

ASCR's Scientific User Facilities Today

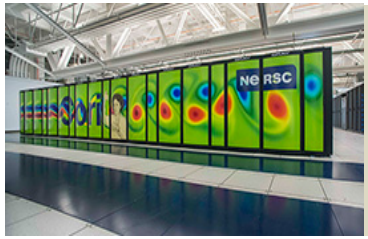

National Energy Research Scientific Computing Center (NERSC)
Upgrade:
Perlmutter Cray/AMD/NVidia

- **High-end capacity computing for the Office of Science research community**
- **Approximately 7,000 users and 800 projects**

Argonne Leadership Computing Facility (ALCF)
Upgrade: Aurora Intel/Cray

Oak Ridge Leadership Computing Facility (OLCF
Frontier: Cray/AMD/AMD

- **Highest computational capability**
- **Open to all**
- **Approximately 1,000 users and 50-60 projects at each center, each year**

Energy Sciences Network (ESnet)
Upgrade: Esnet 6

- **DOE's high-performance research network facility connects all 17 DOE national laboratories plus many additional DOE sites and is tuned for science applications**
- **Enabler for tens of thousands of researchers**

More users and groups are leveraging multiple facilities

The width of the ribbon connecting two facilities corresponds to the number of users who utilized both of those facilities

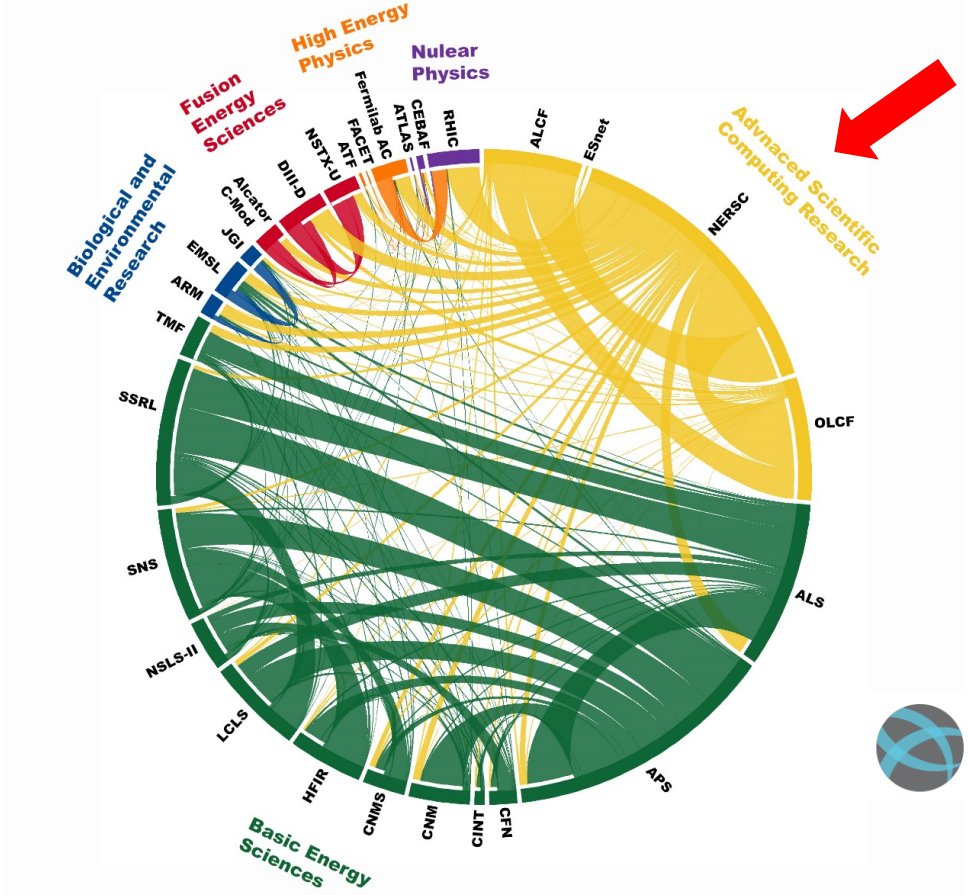

ESnet
ENERGY SCIENCES NETWORK

=

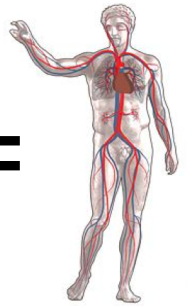

Overview of Workshops

Research priorities for the management, analysis, and visualization of experimental and observational data (EOD). Science user facilities will be soon collectively acquiring exabytes of data per year (2015)

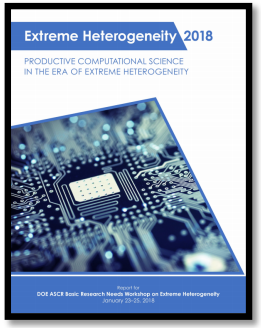

Priority Research Directions for realizing the capabilities needed to address the challenges posed by “extreme heterogeneity”, systems built from a custom aggregation of components; and the difficulty and complexity of developing scientific software. (2018)

Priority Research Directions (PRDs) as viewed through the lens of applied mathematics and scientific computing (2018)

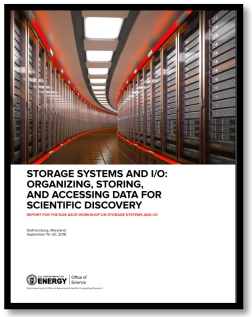

Key challenges and research directions that will advance the field of storage systems and I/O over the next 5–7 years. Addressing these combined challenges and opportunities requires tools and techniques that greatly extend traditional approaches. (2018)

Priority Research Directions to support current and future scientific computing needs, which will increasingly incorporate a number of different tasks that need to be managed along with the main simulation or data analysis tasks. (2019)

Data for AI (2019)

Coming Soon

Priority Research Directions:
Components and Capabilities Needed for Successful ISDM

Pervasive, controllable, composable, and transparent ISDM co-designed with the software stack and with fundamentally new algorithms.

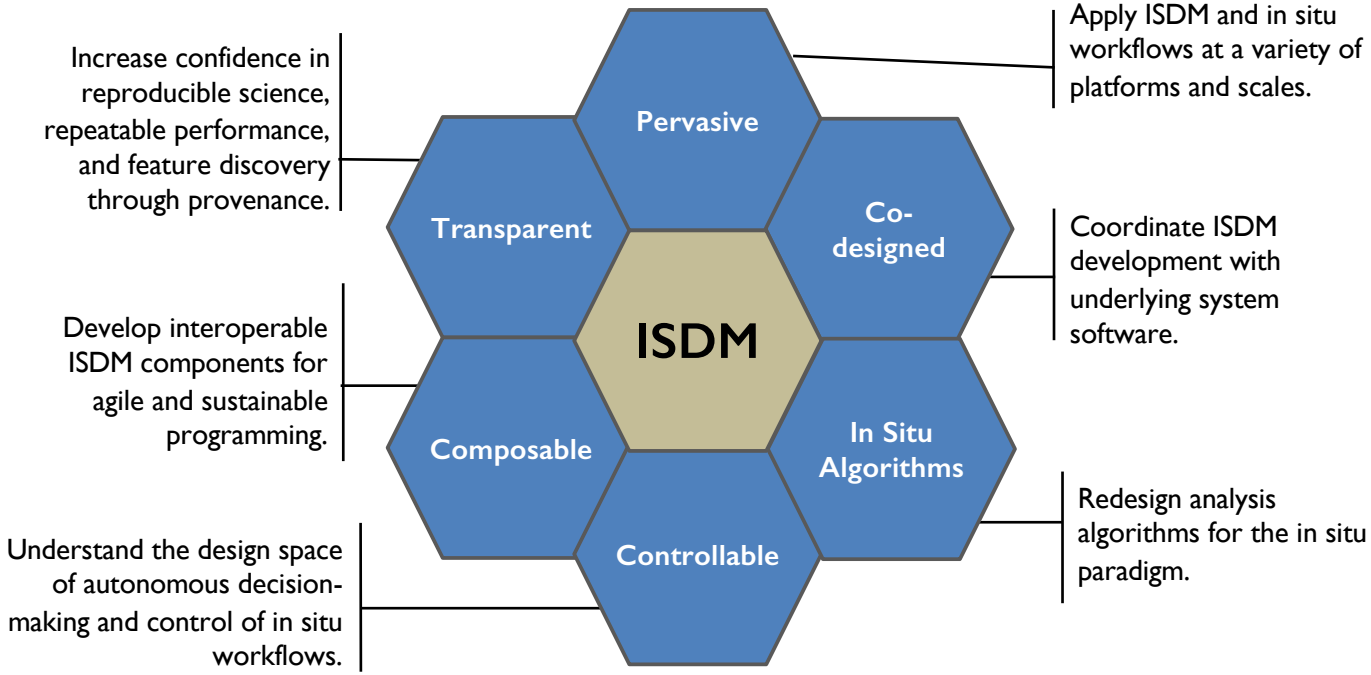

Office of Science Data for AI Round Table:
Challenges, Opportunities, & Enabling Capabilities

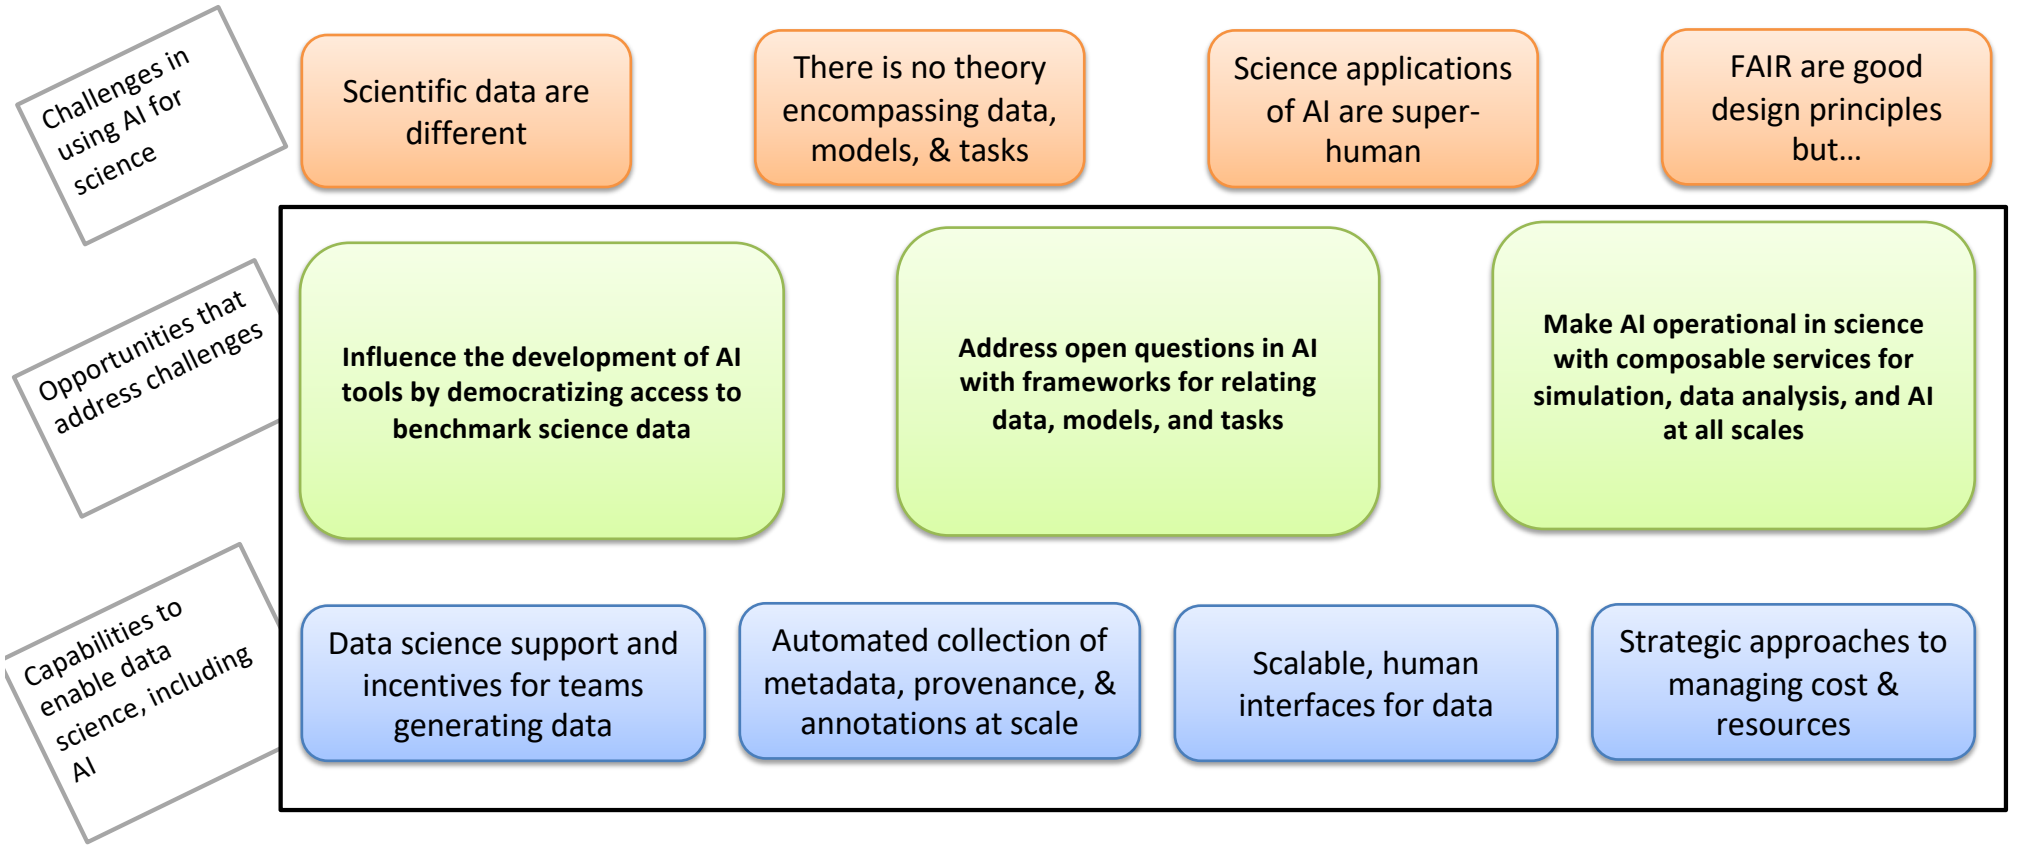

ASCR Exascale Requirements Review Workshop Reports

http://exascaleage.org/

Experimental and Observational Data

Executive Summary

- Gaining scientific knowledge from experimental data is increasingly difficult
- Convergence of data and computing: data- and computing-centric needs increasingly intertwined, symbiotic
- Acute, urgent data-centric needs in SUFs and science programs

Strategic Vision for ASCR Research Program

Emerging trends are pointing to a future that is increasingly

1. **Instrumented**: Massive data, high-tech sensors, detectors, satellites
2. **Interconnected**: Internet of Things, heterogeneous & composable resources
3. **Automated**: Machine learning for complex processes, real-time requirements
4. **Accelerated**: Fast & flexible pathways for open science & research insights

What is the role of the ASCR Research Program in transforming the way we carry out energy & science research?

1. **Post-Moore technologies**: Need basic research in new algorithms, software stacks, and programming tools for quantum and neuromorphic systems
2. **Extreme Heterogeneity**: Need new software stacks, programming models to support the heterogeneous systems of the future
3. **Adaptive Machine Learning, Modeling, & Simulation for Complex Systems**: Need algorithms and tools that supports automated decision making for intelligent operating systems, in situ workflow management, resilient infrastructure, & improved operational capabilities
4. **Uncertainty Quantification:** Need basic research in uncertainty quantification and artificial intelligence to enable statistically and mathematically rigorous foundations for advances in science domain-specific areas

5. **Data Tsunami:** Need to develop the software and coordinated infrastructure to accelerate scientific discovery by addressing challenges and opportunities associated with research data management, analysis, and reuse

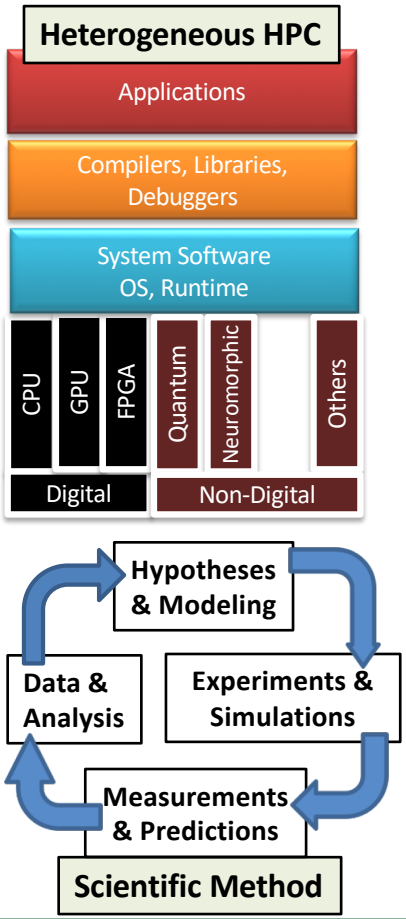

Scientific Machine Learning Priority Research Needs

Scientific Machine Learning: Foundations

Domain-Aware: Leverages & respects scientific domain knowledge. Physics principles, symmetries, constraints, uncertainties & structure-exploiting models

Interpretable: Explainable and understandable results. Model selection, exploiting structure in high-dimensional data, use of uncertainty quantification with machine learning

Robust: Stable, well-posed & reliable formulations. Probabilistic modeling in ML, quantifying well-posedness, reliable hyper-parameter estimation

Scientific Machine Learning: Capabilities

Data-Intensive SciML: Scientific inference & data analysis. ML methods for multimodal data, in situ data analysis & optimally guide data acquisition

Machine Learning-Enhanced Simulations: ML hybrid algorithms & models for predictive scientific computing. ML-enabled adaptive algorithms, parameter tuning & multiscale surrogate models

Intelligent Automation and Decision Support: Adaptivity, automation, resilience, control. Exploration of decision space with ML, ML-based resource management, optimal decisions for complex systems

January 2019

Advances in 6 Priority Research Directions (PRDs) are needed to develop the next generation of machine learning methods and artificial intelligence capabilities.

Upcoming ASCR Workshop on Data Reduction

- **Workshop will inform ASCR data reduction strategic planning**
  - Experimental data and simulation data
  - More than just data compression…
- **Joint ASCR effort:**
  - Applied Mathematics
  - Computer Science
- **Contact information:**
  - Laura Biven: laura.biven@science.doe.gov
  - Bill Spotz: william.spotz@science.doe.gov
  - James Ricci: james.ricci@science.doe.gov
- **Still to be determined:**
  - Dates (April 2020 or later)
  - Organizing committee

Backup Slides

Productive Computational Science in the Era of Extreme Heterogeneity:

Revolutionizing How We Utilize Leadership Class Computers for Scientific Discovery

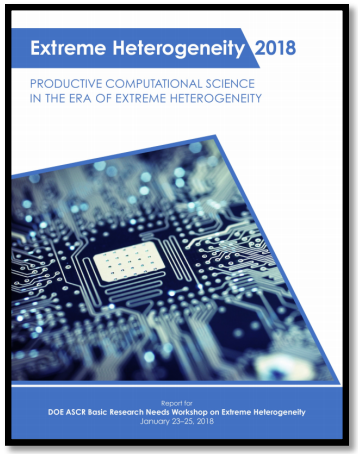

Drivers

- **Scientific**
  - Increasingly diverse science workflows
  - Increasingly diverse science users
- **Technological**
  - **Market-driven innovation**
    - **Data driven workloads**
    - **AI/ML**
  - Market-based constraints
    - Conserve power & heat
    - International competition

Key Questions

- *What will improve productivity of scientific software developers?*
- ***Can AI /ML help coordinate & control diverse computing resources?***
- *Can advanced mod-sim predict performance for application designers & improve system ROI?*
- *What will improve verifiable scientific findings in the future?*
- *What software infrastructure will we need for productive scientific workflow across multiple complex computing environments?*

Priority Research Directions

- **Maintaining and improving programmer productivity**
- **Managing resources intelligently**
- **Modeling & predicting performance**
- **Reproducible science**
- **Facilitating data management, analytics, & workflows**

CAMERA: Center for Advanced Mathematics for Energy Research Applications
A joint ASCR-BES Center for new algorithms & data analysis at DOE user facilities

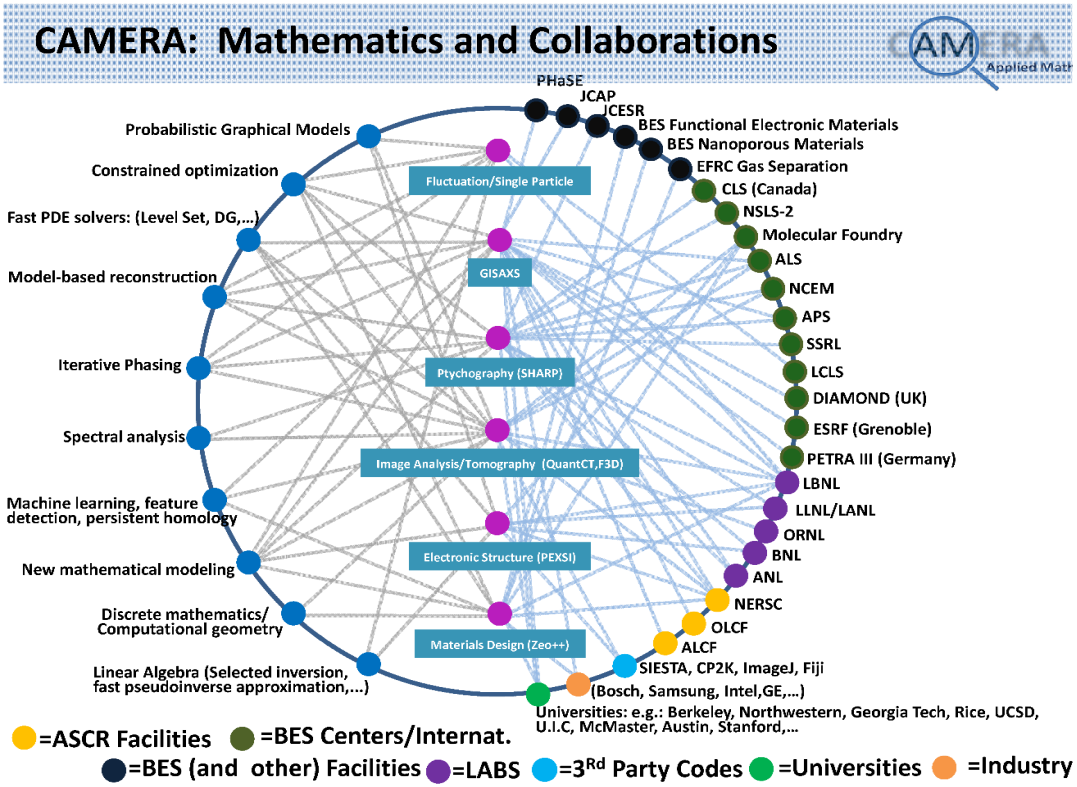

CAMERA research & codes are the centerpiece that links new mathematics, applications, and collaborators.

Connected graph shows that:

- Different types of mathematics (on left) can have an impact on multiple codes & applications
- Mathematics for one particular application can suggest new ideas when addressing different application challenges

Machine learning and AI are emerging technologies that will increase CAMERA's impact on DOE user facility applications.